

Deployment Considerations for Dual-stack Lite

draft-lee-softwire-dslite-deployment-00

Yiu Lee,
Roberta Magione,
Carl Williams,
Christian Jacquenet
Mohamed Boucadair

DS-lite Deployment Considerations

- Based on preliminary experimental deployment, this work describes deployment and operation considerations for DSLITE.

DS-lite Architecture

- It is recommended that the AFTR addressing architecture should consist of two individual interfaces (i.e. one dedicated for IPv4 and one dedicated for IPv6) to segregate the functions.

MTU Considerations

- With DS-lite (as with tunneling protocols) comes additional header overhead that implies that the tunnel's MTU is smaller than the raw interface MTU.
- The issue that the end user will experience is that they cannot download Internet pages or transfer files using File Transfer Protocol (FTP) but may be able to ping successfully.

Lawful Intercept Considerations

- RFC 2804, "IETF Policy on Wiretapping", says that the IETF will NOT work on wiretapping functionality in IETF standards documents. However, the IETF does encourage the publication of wiretapping mechanisms for broad community review.

Lawful Intercept Considerations

- Interception in DS-lite architecture could be performed on the AFTR itself.
- Time-stamped logging of the address and port mappings at the AFTR should be maintained, which in turn can add resource burden to the AFTR devices.

Logging @ AFTR

- The time-stamped logging is also important for tracing back specific users when a problem is identified from the outside of the AFTR.
- Policies applying to incoming sources must be implemented on the outside of the AFTR. Once the packets are translated, they cannot be easily identified by IPv4 address without some correlation with the AFTR mapping table.

AFTR Policies

- Policies applying on the NAT-ed addresses should be implemented on the external interface of the AFTR.
 - Once the packets are translated, they cannot be easily identified by IPv4 address without some correlation with the AFTR mapping table.
- Policies applying to outgoing sources should be implemented on the customer-facing side of the AFTR for the same reason.
 - In order to be able to deploy different services offers, multiple set of policies can be configured on the AFTR: each set of policies can then be applied to a different logical tunnel interface on the AFTR.

AFTR Impacts on Accounting Process in Broadband Access

- The accounting process at the AFTR level is only necessary if the Service Provider requires separate per user accounting records for IPv4 and IPv6 traffic.
- If the per user IPv6 accounting records, collected by the BNAS, are sufficient, the additional complexity to be able to implement IPv4 accounting at the AFTR level is not required.

Reliability Considerations of AFTR (1/3)

- The service provider can use several techniques to achieve high availability such as various types of clusters to ensure availability of the IPv4 service.
- DS-lite HA techniques include **cold standby mode**:
 - When the Primary AFTR fails, all the existing established sessions will be flushed out. The internal hosts are required to re-establish sessions to the external hosts.

Reliability Considerations of AFTR (2/3)

- **DS-lite Hot standby mode:**
 - AFTR keeps established sessions while failover happens. AFTR states are replicated from the Primary AFTR to the Backup AFTR. When the Primary AFTR fails, the Backup AFTR will take over all the existing established sessions.
 - In the DS-lite Hot standby mode, the internal hosts are not required to re-establish sessions to the external hosts.

Reliability Considerations of AFTR (3/3)

- **Combo mode** is a method to deploy DS-lite between these two whereby only selected sessions such as critical protocols are replicated.
 - Criteria for sessions to be replicated on the backup would be explicitly configured on the AFTR devices of a redundancy group.

Placement of AFTR (1/2)

- The AFTR architecture design is the strategic placement of each AFTR to best use the capacity of each public IPv4 address without oversubscribing the address or overtaxing the AFTR itself.

Placement of AFTR (2/2)

- It is important to centralize the public IPv4 addresses where each address no longer represents a single machine, a single household, or a single small office.
- The address now represents multiple machines, homes, and offices related only in that they are behind the same AFTR.
- An issue of the placement of AFTR is the identification by IP address as it becomes difficult and thus applications that assume such geographic information may not work as intended.

Geographic aware applications...

- It is important to locate the AFTR so that various applications and services will place their servers in such a way to locate them near sets of user so that this will lessen the latency on the client end.
- Having sufficient geographical coverage can indirectly improve end-to-end latency.

DS-lite impacts on QoS

- As with tunneling in general there are challenges with deep packet inspection with DS-Lite for purposes of QoS.
- Service Providers commonly uses DSCP to classify and prioritize packets.
 - It is recommended the AFTR and B4 should copy the DSCP value in the IPv4 header to the IPv6 header after the encapsulation.

Port Forwarding Considerations

- Some applications require accepting incoming UDP or TCP traffic.
- Some applications rely on ALGs, UPnP IGD, or manual port configuration. Port Control Protocol (PCP) [I-D.wing-pcp-design-considerations] is designed to address this issues.

B4 Deployment Considerations

- In order to configure the IPv4-in-IPv6 tunnel, the B4 element needs the IPv6 address of the AFTR element.
 - This IPv6 address can be configured via an out-of-band mechanism, manual configuration or a variety of DHCPv6 options.
- It is recommended that in order to have interoperability that the B4 element should implement the DHCPv6 option defined in [I-D.ietf-softwire-ds-lite-tunnel-option].

DNS Proxy @ B4

- B4 should contain a DNS proxy resolver and forward DNS queries to an external recursive resolver over IPv6.
- Alternately, the B4 proxy resolver can be statically configured with the IPv4 address of an external recursive resolver.
 - Here the DNS traffic to the external resolver will be tunneled through IPv6 to the AFTR which will consume NAT resources (NOT Recommended)

Security issues

- Some of the security issues with carrier-grade NAT result directly from the sharing of the routable IPv4 address apply with DS-lite.
 - I.E... Devices on the customers side may try to carry out general attacks against systems on the global Internet or against other customers by using inappropriate IPv4 source addresses inside tunneled traffic.
- In short, the AFTR entity must protect against such attacks.

Summary

- Deployment considerations of the B4, AFTR and DNS have been discussed and recommendations for their usage have been discussed.
- It is the goal that this document and discussion can be a reference for service providers and network providers deploying DS-lite.